SOLUTION OVERVIEW

MCP Gateway for Agents

Identity, Containment, and Audit for Autonomous AI

PRODUCT BRIEF FOR AKAMAI CONNECTED CLOUD

When the Caller Isn't Human

AI agents don't just answer questions anymore — they act. They run headless, chain tool calls through Model Context Protocol (MCP), retry aggressively, and operate on authority delegated by the people they work for. Governance built around human logins doesn't hold when no one is watching each call. MCP Gateway treats agents as first-class principals, with their own identity, scoped authority, and hard limits.

Why Agents Need a Gateway

- **Their Own Identity**
  Every agent must be a known principal — registered, authenticated, and named in every log line — not an anonymous process hiding behind a shared account.
- **Provable Delegation**
  When an agent acts for a person, the tool it calls should know exactly who that person is — with cryptographic proof, not a super-token everyone shares.
- **Least Privilege, Enforced**
  An agent should see exactly the tools its job requires — and nothing else. Access must default to deny and fail closed.
- **Hard Ceilings**
  Agentic loops go wrong fast: retry storms, recursive calls, sudden fan-out. Containment can't depend on the agent behaving.

Inside AAM: MCP Gateway for Agents

Akamai AI and API Manager (AAM) gives every agent a real identity and wraps every tool call in enforcement — on Akamai Connected Cloud.

- **Agents as OAuth Clients**
  Dynamic client registration with PKCE gives each agent its own identity. Headless agents authenticate with ID-JAG assertions, cloud workload identity, mTLS, or API keys — no browser required.
- **Delegation with ID-JAG**
  Cross-App Access token exchange carries "acting for Alice" cryptographically: short-lived, single-audience assertions with fine-grained authorization constraints.
- **Least Privilege Per Agent**
  Default-deny catalogs scope each agent to the tools its job requires. Non-permitted tools are hidden from discovery and refused on invocation — fail-closed.
- **Runaway Containment**
  Per-agent rate limits, quotas, timeouts, and response-size caps hold no matter what the agent decides to do. Spikes are contained at the edge in real time.
- **Fail-Closed Guardrails**
  No human sanity-checks each call, so the gateway does: schema validation rejects malformed calls, prompt-injection and tool-poisoning screening guards the agent's context, and DLP stops exfiltration.
- **Every Action Attributed**
  A typed event stream records which agent, on whose behalf, called which tool, with what outcome — with OpenTelemetry tracing and real-time SIEM export.

AI and API Manager
Built into Akamai Connected Cloud

SOC 2 Type 2 | OAuth 2.1 | ID-JAG | OpenTelemetry | Prompt-injection protection